https://doi.org/10.2352/CIC.2022.30.1.29

The Regularised Epsilon-Derivative for Image Reintegration

Graham D. Finlayson[1] and Mark S. Drew[2]
[1] **School of Computing Sciences, University of East Anglia, Norwich, NR4 7TJ, U.K. graham@cmp.uea.ac.uk**
[2] **School of Computing Science, Simon Fraser University, Vancouver, British Columbia, Canada V5A 1S6, mark@cs.sfu.ca**

Abstract

Many image-editing tasks are carried out in the gradient domain. Suppose that for an image I the gradient ∇I consists of a pair of fields (p,q); then some image "reintegration" scheme is tasked with converting derivative fields (p,q) back to image-space I; typically, a Poisson equation solver is used for this task. But what if we have altered (p,q) so that this pair (p,q) is no longer integrable? Then we have to project onto integrable gradient data that will indeed reintegrate to an approximation of the original image. For example, we may wish to alter (p,q) so as to emphasize or de-emphasize some aspects of the image, e.g. ameliorating wrinkles in skin images (or indeed enhancing them in the case of ageing a face image).

Here, we propose a new gradient kernel that retains part of the original image, regularising the reintegration back into the image domain. We compare our approach with the Screened-Poisson method which includes a term λ times a "screen" term that moves the solution image back closer to the input image. Effectively, we are doing a similar adjustment, but we show that the results are a good deal better than using Screened-Poisson, which tends to overly blur the output. Moreover, in Screened-Poisson one must choose a value for λ, which may be different for every image – here we determine that our new kernel method does not need to adapt to each image yet delivers comparable or better results.

1. Introduction

Gradient domain processing has attracted significant interest due in part to the importance of edges in human perception [1]. And indeed many image-editing tasks are carried out in the gradient domain. For example, Poisson Image Editing [2] involves making changes in the gradient field, and then reintegrating back to image space. But the reintegration step requires some care. Indeed, the altered gradient field is no longer an integrable pair (p,q) and we have to project the approximate gradient onto an integrable pair by using some variational method such as a Poisson equation.

To give some intuition about 'non integrability', an edge (or gradient) in a 2-d image has two components: a derivative in the x- and a derivative in the y-direction. If we "made up" per pixel x- and y-gradients there is no guarantee that there is actually an image (which only has one intensity per pixel) that has these gradients. Reintegrating by solving the Poission equation finds the image whose gradients are closest to the desired ones in a least-squares sense.

In this paper, we propose that the gradient can be augmented by making use of a new gradient kernel that retains part of the original image itself. This regularizes the projection by utilising the extra part, carrying information about the image itself. In effect, the augmented gradient, by construction, results in a more integrable field (and solving for the most consistent reintegration is more straightforward). Our augmentation idea is straightforward to implement. In the discrete image domain a derivative (in the x-direction) is computed as $I(x+1,y)-I(x,y)$. Under our augmented scheme, the ε-derivative is computed as $(1+\varepsilon)\cdot I(x+1,y)-I(x,y)$. The y epsilon derivative is computed analogously.

Our idea is reminiscent of the Screened-Poisson method, which augments the variational objective functional with a term λ times a "screen" term which typically consists of the image itself. An issue in that method is that we need to decide upon a value for λ. Here we compare our new kernel method with Screened-Poisson and find that while λ in the latter method needs to adapt for each image, the kernel for our new method can be chosen once and for all and be applied to any input image.

To evaluate our method and compare to the Screened-Poisson approach we need to choose an gradient domain application. There are, in fact, a plethora of gradient-domain methods that could be studied, such as shadow-removal, lightness computation, high-dynamic-range, etc., but here we consider one such problem: we look at the problem of image smoothing by removing high derivatives. The task at hand is to remove a good deal of the gradient information and then reintegrate back to the image domain. The question becomes "what can we recover if we zero out much of the gradient information?". The results show that the new method is *better* than or comparable to Screened-Poisson, but it simpler because we do not need to set any adaptive parameter.

The paper is organised as follows: In §2, we discuss the domain in which we are operating, *viz.* changing the set of image gradients so as to emphasize certain aspects of the image. In §3 we examine the effect of our new scheme when applied to such a problem. Section 4 presents results on an extensive standard set of images, and §5 discusses conclusions and future work.

Gradient Manipulations and the ε-Derivative

Many schemas for manipulating images center upon the notion of altering image derivatives to serve some purpose. In an exemplar situation, suppose we can identify shadow-boundaries. Zeroing gradients across such boundaries will tend to produce a reintegrated image with the lighting change attenuated [3, 4]. A substantial collection of gradient-manipulation methods appears in [5].

A feature of such methods is the need to ascend back into the image domain from the gradient domain. Consider one colour-channel image I selected from R,G,B one at a time, and for brevity define $\partial_x I$, $\partial_y I$ to stand for $\partial I/\partial_x$ and $\partial I/\partial_y$. For the image itself, the pair of fields $(p,q)=\nabla I$; $p=\partial_x I$, $q=\partial_y I$ is indeed integrable and we can restore I from (p,q).

Standard Kernel

When (p,q) is not an integrable pair, the usual approach to such a problem leads to a Poisson equation [6]: given a pair of fields (p,q), and for an unknown I, we wish to form a best least-squares solution to the variational problem

$$\hat{I} = \overset{\text{argmin}}{I} \sum \left\{ \|p - \partial_x I\|^2 + \|q - \partial_y I\|^2 \right\} , \tag{1}$$

where $\hat{I}$ is the optimum reconstructed image; this results in a Poisson equation:

$$\partial^2 I/\partial_x^2 + \partial^2 I/\partial_y^2 = \partial_x p + \partial_y q \tag{2}$$

Mathematically, eq.(2) states that the divergence of the gradient for the least-squares solution for I is given by the forcing function on the right hand side.

The solution is best constrained using Neumann boundary conditions [7]. If we use homogeneous Neumann boundary conditions $\partial I/\partial n = 0$ then we remove from consideration extraneous harmonic solutions to Laplace's equation and have a unique solution up to an unknown constant of integration [7].

So far, we have a classical Poisson equation which can be solved iteratively. A non-iterative approach solves the equation analytically in the Fourier domain [8]. For suppose we go over to the Fourier values for both p and q, defining

$$P = \mathscr{F}(p) \, , \; Q = \mathscr{F}(q) \, , \text{ and also } F = \mathscr{F}(I) \tag{3}$$

where $\mathscr{F}$ denotes the Fourier transform.

Now suppose the effect of a derivative filter, when expressed in the Fourier domain is filter a_x for partial derivative ∂_x and a_y for partial derivative ∂_y. Here, let us adopt forward differencing to express partial derivatives. That is, let us use the filter ϕ = $\boxed{-1}\,\boxed{+1}$ for the x-derivative, meaning (-1) times the current pixel plus $(+1)$ times the pixel to the right of the current pixel; and let ϕ^T express the y-derivative (with ^{T} meaning transpose). Then for angular frequency ω_x, ω_y, the filter a_x becomes $\exp(i\omega_x) - 1$ in the Fourier domain. (This filter differs from that in [8] because the latter uses central-differencing as opposed to forward-differencing as here.) Using a 1-based spatial indexing schema for retinal coordinates x, y, if the image resolution is (M,N) then $x = 1..N$, $y = 1..M$, and $\omega_x = 2\pi(x-1)/N$ and $\omega_y = 2\pi(y-1)/M$ (with ω_x and ω_y being $M \times N$ arrays).

Now we can write our functional equation in Fourier coordinates; we have that the best solution for the Fourier transform of the solution image $\hat{F}$ is

$$\hat{F} = \overset{\text{argmin}}{F} \sum \left\{ \|P - a_x F\|^2 + \|Q - a_y F\|^2 \right\} \tag{4}$$

Recall that P and Q could consist of the Fourier domain versions of altered, manipulated gradient values (p,q).

Solving eq. (4) for $\hat{F}$, we arrive at

$$\hat{F} = \frac{a_x^* P + a_y^* Q}{a_x^2 + a_y^2} \tag{5}$$

similar to that in [8] but for forward-differences. Let us call this method the FP (Frankot-Chellappa) solution. Examples of the output for the FC method in various cases are examined below in §3, and compared to our new method.

Screened Poisson

The Screened Poisson (SP) method [9] improves upon the FC method by including a subsidiary term in the objective function (4) meant to drive the solution closer to a "screen" field, typically the original image I itself, to include a data term that the reconstructed function must also approximate. The SP method includes a user-defined scalar parameter λ that expresses the amount of "screening". In this case the variational objective (1) becomes

$$\hat{I} = \overset{\text{argmin}}{I} \sum \left\{ \|p - \partial_x I\|^2 + \|q - \partial_y I\|^2 \right\} + \lambda (I - u)^2 \tag{6}$$

where typically the data term u is the original image itself.

The Euler-Lagrange equation for this modified variational problem is the screened Poisson equation

$$\partial^2 I/\partial_x^2 + \partial^2 I/\partial_y^2 = \partial_x p + \partial_y q + \lambda\,(I - u) \tag{7}$$

A direct method for this is to carry out the needed projection in the Fourier domain. The minimization using the direct method in Fourier space then becomes

$$\hat{F} = \overset{\text{argmin}}{F} \sum \left\{ \|P - a_x F\|^2 + \|Q - a_y F\|^2 \right\} + \lambda\,(F - U)^2 \tag{8}$$

where $U = \mathscr{F}(u)$, with solution

$$\hat{F} = \frac{a_x^* P + a_y^* Q + \lambda U}{a_x^2 + a_y^2 + \lambda} \tag{9}$$

Let us call this method the SP solution. For $\lambda \to 0$ the SP solution moves to the FC solution; and as λ becomes large, the SP solution goes over to the screening function u.

Robust Kernel

As is known, and as we ourselves discover in the experimental section, the best value of the parameter λ in the SP method varies from image to image. Therefore here we set out an alternative approach where a single value of a different parameter is held constant. To do so, we introduce the new kernel ϕ_D = $\boxed{-1}\,\boxed{1+\varepsilon}$. Immediately, we can see that the new operator keeps a partial amount of its argument, as opposed to the familiar ϕ, which consists of the derivative only.

Let us denote the ε-derivative as D_x, D_y. It is helpful to separate the ε-derivative into 2 parts:

$$\begin{aligned} D_x I &= \partial_x I + \chi_x \\ D_y I &= \partial_y I + \chi_y \end{aligned} \tag{10}$$

That is, the usual gradient plus an "extra" two parts χ_x, χ_y:

$$\begin{aligned} \chi_x &= \varepsilon\, I^{\leftarrow} \\ \chi_y &= \varepsilon\, I^{\uparrow} \end{aligned}$$

where $I^{\leftarrow}$ is the image, shifted left; and $I^{\uparrow}$ is the image shifted up. (11)

We refer to the robust kernel operator as an "ε-derivative" [10]. We argue in [10], and show there an embodiment in terms of the HDR problem, that due to the regularization brought about by use of the ε-derivative, halos and other image artifacts can be ameliorated. The new operator provides a more robust derivative structure since it retains part of the image information lost in the standard gradient.

Hence, the method proposed is as follows. Suppose we wish to alter the gradient of an image in some meaningful way — for example, here as a test case we set to zero all high-magnitude derivatives, thus removing high-frequency features and effectively smoothing the image. The most simple approach for the smoothing problem as stated is to form the gradient ∇I for an input image I and, using a Boolean mask aimed at removing high values of the gradient, set to zero targeted gradients by multiplying by the mask, thus forming the pair (p,q), an approximation of a gradient.

Recall that the new operators D_x, D_y consist of the gradient plus an extra term. So we propose here to indeed threshold the image gradient $\partial_x I, \partial_y I$ part of $D_x I, D_y I$ but keep the extra terms stemming from the ε-derivative from the original image unchanged, so that part of the original image persists in the new operator, as we desired.

Now we again go over to Fourier space. As an extension of the FC solution, now we let

$$\begin{aligned} P &= \mathscr{F}(p + \chi_x) \\ Q &= \mathscr{F}(q + \chi_y) \end{aligned} \tag{12}$$

where χ_x, χ_y are the extra parts of D_x, D_y.

We have the same objective function as in eq. (4) and indeed the same solution eq. (5) as in the FC solution; but with the crucial difference that now we are operating with the full D_x, D_y operators and full extended parts: the filters are now the robust Fourier-filters $a_x = \exp(i\omega_x) - (1+\varepsilon)$, $a_y = \exp(i\omega_y) - (1+\varepsilon)$.

3. Assessment of ε-Derivative

Let us consider a standard gradient-manipulation scenario wherein we smooth an image, by zeroing large-value gradients. Consider the image Fig. 1(a). Here we mean to use an exemplar image in the first instance, in order to guide the development and understand all that takes place for a suitable image.

To provide an aggressive operation, so as to be able to clearly see the effects of the various solutions, suppose we remove all gradients that are over the 50-percentile of gradient magnitude. Define the gradient magnitude as

$$mag \;=\; \left(\sum_{k=1}^{3} (\partial_x I_k)^2 \;+\; \sum_{k=1}^{3} (\partial_y I_k)^2 \right)^{1/2} \tag{13}$$

and threshold

$$\begin{aligned} mask \;&=\; \left(mag < 50^{\text{th}} \text{ quantile of } mag \right) \\ \nabla' I \;&=\; \nabla I \times mask \end{aligned} \tag{14}$$

This says that gradients greater than the 50th quantile are set to zero.

As a first effort, consider Fig. 1(b), which shows the result of standard Gaussian smoothing on the input image. This is a useful result, but arguably too blurred. We would like to have a smoothed result, but one that retains a good deal of the information in the image. Moreover, we wish to test the ε-derivative by zeroing out some of the gradient information.

So to continue, let us zero out high-gradient values: Fig. 1(c) displays the pixels where we do not zero out the gradient (i.e., where the mask is 1); that is, where we retain low-magnitude gradients. In the present application this is given by 50% of the gradients, of course, and the mask image shows how substantial a removal of the gradient information we are carrying out.

Then the FC solution is as is shown in Fig. 2(a). This is clearly a result that is not of much use to us: removing such a substantial amount of gradient information does not generate an image adhering to either the colour or texture of the original. That the result is so poor is an indication that the gradient field is far from being integrable.

However, looking at the original in Fig. 1(a), re-rendering this image so that the large derivatives are set to 0, our expectation is broadly that the output image should look like the input but with any large changes eliminated, with resulting smoothing.

Now, we consider the ε-derivative method, we arrive at Fig. 2(c). We have found that we can in fact hold fixed the ε parameter in this method, using $\varepsilon = 0.2$. Clearly, the result gels well with our expectations. The image looks smoothed but because many small derivatives are kept there is good image detail too.

Now, let us compare against the the SP method.Here we have to set the adaptive parameter λ; to favour the SP method as much as possible, we quantify how closely the SP solution for a particular λ matches the ε-derivative solution below, in terms of the PSNR (peak signal to noise ratio). We traverse λ values in the range $[0.0, 0.2]$, where images are scaled to values $[0.0, 1.0]$. A typical best-result occurs with λ between 0:01 and 0.10. Using the value 0.01, we obtain the result in Fig. 2(b), a better result than for the FC solution because the screened method pulls the solution back towards the input image; but the result is so smoothed its utility is reduced. Clearly, in terms of this example the ε-derivative method delivers the *best* smoothed output image.

To quantify the relationship between these smoothed images and the original image, we need a metric $\delta(I, I')$ that takes into account both colour, as in using PSNR over the 3 colour-channels, and also texture, since we are smoothing images. Here we suggest the Kullback-Leibler Divergence distance (the relative entropy) from the original image, but calculated just for the $(mask = 0)$ pixels, i.e., how well does an algorithm regenerate the missing-gradient information, as shown in the reintegrated image.

For the SP method, and using 128-bin histograms, we find a KL divergence of 0.0426 bits. Whereas for the Epsilon-derivative solution we find a KL divergence of 0.0124 bits, or about 4 times better.

4. Further Tests

In Fig. 4, we run the competing methods on the 24 images in the Kodak-CD image database.[1] We see that, just as for our test image, the FC solution is not useful, the screened SP solution is reasonable but blurry, and the Epsilon-derivative solution is the best output. For the 24 images, some of which are shown in Fig. 4, Fig. 3 shows a histogram of the value of each best-value of λ,

[1] http://r0k.us/graphics/kodak/

one for each image, such that the SP solution best matches the epsilon-derivative result, the measure being the PSNR for the SP solution compared to the epsilon-derivative result. We see that the SP method may work comparably to the epsilon-derivative method, but often does not and cannot be relied upon.

5. Conclusion

In this work we have presented a novel, alternative Fourier-domain approach for solving for a projection of a set of nonintegrable fields (p, q) to a closest reintegrated image. The new idea is to use a robust form of gradient, the ε-derivative, to regularize reintegrating from the gradient domain back to the image domain. We found in tests that we could settle on values $\varepsilon = 0.2$ for any image.

Future work includes expanding the set of problems investigated in terms of the new ε-derivative operator.

References

[1] T. Cornsweet. *Visual Perception*. Academic Press, 1970.

[2] P. Pérez, M. Gangnet, and A. Blake. Poisson image editing. *ACM Trans. Graph.*, 22(3):313–318, 2003.

[3] Stuart E. Lynch, Mark S. Drew, and Graham D. Finlayson. Colour constancy from both sides of the shadow edge. In *CPCV2013: Int. Conf. on Computer Vision Workshop on Color and Photometry in Computer Vision*, Dec. 2013.

[4] G.D. Finlayson, M.S. Drew, and C. Lu. Entropy minimization for shadow removal. *Int. J. Computer Vision*, 85(1, Oct):35–57, 2009.

[5] Pravin Bhat, C. Lawrence Zitnick, Michael Cohen, and Brian Curless. Gradientshop: A gradient-domain optimization framework for image and video filtering. *ACM Trans. Graph.*, 29(2):10:1–10:14, 2010.

[6] B. K. P. Horn. *Robot Vision*. MIT Press, 1986.

[7] A. Blake. Boundary conditions for lightness computation in Mondrian world. *Comp. Vision, Graphics, and Image Proc.*, 32:314–327, 1985.

[8] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 10:439–451, 1988.

[9] P. Bhat, B. Curless, M. Cohen, and L. Zitnick. Fourier analysis of the 2D screened Poisson equation for gradient domain problems. In *ECCV08*, 2008.

[10] Graham D. Finlayson and Mark S. Drew. U.S. Patent (Apple corp.): Regularized derivative operators for image processing system and method, Granted as U.S. Patent 11,238,573, Feb. 1, 2022.

(a) (b) (c)

Figure 1. *Smoothing: Suppress large gradients above 50-percentile. (a): Input image. (b): Gaussian smoothing. (c): Mask for retained gradients.*

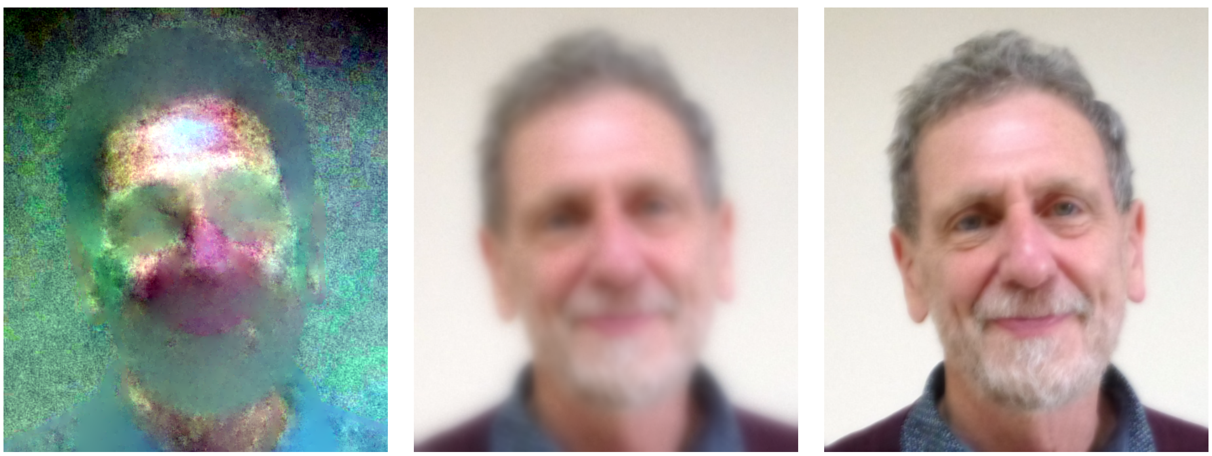

(a) (b) (c)

Figure 2. *Smoothing: (a): FC solution. (b): SP solution. (c): Epsilon-derivative solution.*

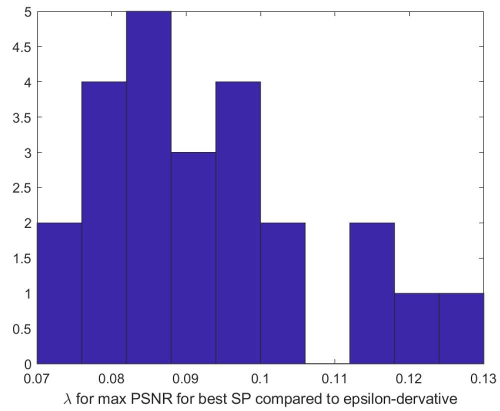

Figure 3. *For each image of the 24 in the database, favour the SP method as much as possible by choosing λ to match the epsilon-derivative solution as closely as possible.*

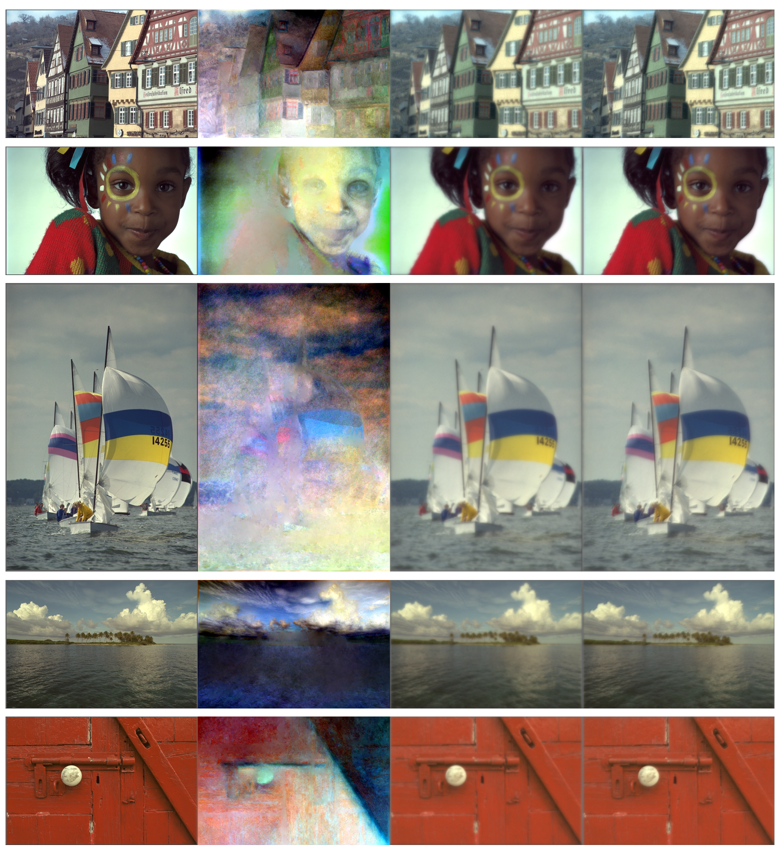

Figure 4. *Kodak-CD database of 24 images. Smoothing: Input image, FC solution; SP solution; Epsilon-derivative solution. Images shown are those with [min, 5-percentile, median, 95-percentile,max] PSNR of SP method, compared to the epsilon-derivative method. The λ values for these are $\lambda =$[0.080, 0.070, 0.120, 0.090, 0.095], for PSNR values [26.807, 29.573, 31.729, 35.874, 36.318] dB.*